DOCUMENT RESUME

ED 212 976 CS 006 484

AUTHOR Kay, Paul
TITLE Three Properties of the Ideal Reader.
PUB DATE Mar 82
NOTE 34p.

EDRS PRICE MF01/PC02 Plus Postage.
DESCRIPTORS *Cognitive Processes; Discourse Analysis; *Expectation; Language Processing; *Learning Theories; Prior Learning; *Reading Processes; *Reading Research
IDENTIFIERS Inferences; *Reading Strategies; *Schemata

ABSTRACT

The main experience of an ideal reader while reading a text is an "envisionment" of that text, a representation in the reader's mind of the content of the text. According to this view the envisionment grows and sometimes changes as the reader progresses through the text, and the ideal reader not only updates and supplements the envisionment of the world the text is describing but also formulates hypotheses, asks questions, notes evidence, and generally accomplishes a variety of processing operations. Based on a variety of examples, there are three properties that an ideal reader uses to construct the envisionment of a text, two of which have to do with the structure of the envisionment, while the third concerns the embedding of both the envisionment itself and the processing operations that produce it. The first property is the ability to "read between the lines," to draw inferences that are not directly warranted by the grammar and lexicon of the text but that are still necessary to derive a coherent envisionment. The second property is the incorporation of the idea of a trusting reader into the concept of an ideal, sophisticated reader; thus, an ideal reader becomes aware of the expectations one has about reading and how these expectations can be manipulated by the author and the text. The third property involves a parsimony principle; that is, the organization of the envisionment--its expectations (schemata), its inferences, and its resulting hypotheses--is arranged in the simplest scenario possible. (RL)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED212976

U.S. DEPARTMENT OF EDUCATION
NATIONAL INSTITUTE OF EDUCATION
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☐ This document has been reproduced as received from the person or organization originating it.

☑ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official NIE position or policy.

Three Properties of the Ideal Reader

Paul Kay

March 1982

"PERMISSION TO REPRODUCE THIS MATERIAL HAS BEEN GRANTED BY

Paul Kay

TO THE EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)."

CS006484

Introduction: The Ideal Reader

One way to study the acquisition of the ability to read is to compare the fledgling reader's attempts at reading with the performance of a hypothesized ideal reader. An approach of this kind has been developed by a research project of which I am a part and on whose collective work I report here.[1] The present essay does not attempt a full report of our project to date, nor even a full characterization of the concept 'ideal reader' (for which see Fillmore 1981). Rather I sketch only the bare outlines of the concept and then concentrate on three important properties of the ideal reader.

The ideal reader is defined in terms of a particular text and a particular interpretation of that text. Given the text and an interpretation, the ideal reader is a device that is possessed of just the knowledge and skills required to extract that interpretation from that text. The ideal reader knows what the text presupposes and is able to learn what the text is designed to convey. The ideal reader thus conceived is not to be confused with the mature, competent reader; rather the ideal reader may be instantiated by a mature, competent reader of A WELL / CONSTRUCTED TEXT WITH A STRAIGHTFORWARD INTERPRETATION. The distinction is illustrated by text (1), with the intended interpretation given in 'single quotes'.

(1) When Pat and Leslie met she smiled but the other didn't smile back.

'When they met, Pat smiled but Leslie didn't smile back.'

Since the ideal reader is defined with respect to a text and an interpretation, the ideal reader of (1) will by definition achieve the interpretation given. But of course mature, competent readers (or hearers)

will not unambiguously assign this interpretation to text (1). Thus, as well as providing a standard against which the performances of real readers can be judged, the concept of 'ideal reader' gives us a basis for criticizing texts. If a text makes unrealistic or undesirable demands on its reader -- e.g., that the ideal reader for this text be clairvoyant or that it be content to believe contradictions: demands that we think are not made on actual, mature and competent readers by well constructed texts -- we judge the text on that account to be deficient.[2]

The main thing an ideal reader does as it reads a text is construct an 'envisionment' of that text, a term we have borrowed from John Seely Brown. The envisionment is the representation in the reader's mind of the content of the text. The envisionment grows, and sometimes changes, as the reader (or hearer of a monologue) progresses through the text.[3] According to this view, the ideal reader, as it reads a text, not only updates and supplements its envisionment of the world the text is describing, but also -- in the service of building this envisionment -- formulates hypotheses, asks questions, notes evidence and in general, accomplishes a variety of processing operations (Fillmore and Kay 1980: 22-49). These processing operations will not be discussed here. Rather, two of the properties of the ideal reader that I shall discuss have to do with the structure of the envisionment and the third concerns the embedding of both the envisionment itself and the processing operations that directly produce it in higher order processes. Instead of attempting further abstract characterization of these properties of the ideal reader, I will proceed directly to examples that illustrate the notions I have in mind.

Levels of Envisionment

The first property of the ideal reader is that it does not invest in every part of its envisionment the same degree of confidence. The ideal reader does a lot of reading between the lines. But the ideal reader also knows that it must place less trust in those parts of its envisionment that arise from reading between the lines than in those parts that come from reading the lines themselves. The ideal reader in fact distinguishes many levels of confidence within the envisionment. Consider the two text fragments:

(2) The Orioles' shortstop bunted the ball right to the first baseman, who grabbed it and tagged the batter out.

(3) The Orioles' shortstop threw the ball right to the first baseman, who grabbed it and tagged the batter out.

Let us catologue several inferences the ideal reader will draw from these texts, that is, several items that ought to become part of the ideal reader's envisionment of these texts. In particular we want to look at relations of coreference between pairs of noun phrases in the two texts. First of all, we note that in both (2) and (3) the pairs of noun phrases (the ball, it) and (the first baseman, who) are coreferential, which facts are summarized in Table 1.

	Text (2)		Text (3)
(a)	the ball = it	(a)	the ball = it
(b)	the first baseman = who	(b)	the first baseman = who

Table 1

Facts (a) and (b) reflect directly the grammar of the texts, and nothing more. This may be seen by considering text (4), which is parallel in grammatical structure to both texts (2) and (3) (which themselves are

grammatically parallel, differing only in the alternation between the lexical items bunt and throw), but in which nonsense words are substituted for the content items.

(4) The Wimbats' glurb slunked the wint to the girfman, who critched it and ...

The reader will appreciate that the grammar of (4) determines the relations of coreference must be as shown in (5), which accord with those of Table 1.

(5) The Wimbats' glurb_i slunked the wint_j to the girfman_k, who_k critched it_j, and ...

Thus inferences (a) and (b) are not dependent on any knowledge or skill the reader may possess beyond his or her knowledge of English grammar. In particular, no knowledge of the game of baseball is required. Also, and more importantly, it is impossible to imagine a continuation of text (2) or (3) that could overthrow these inferences. For example, if we load the content to make conflicting inferences regarding coreferences more plausible, the result is not a switch of perceived relations of coreference but a judgement of incoherence. Thus in (6) we rig the content to try to get the relative pronoun who to be coreferential, not with the immediately preceding noun phrase, but with the initial noun phrase; the attempt fails.

(6) *The batter_i grounded to the shortstop_j, who_i was thrown out easily.

In short, the reader need know nothing about baseball and, in particular, need know nothing about the meanings of the words shortstop, bunt, first baseman, and so on, to infer the facts given in Table 1. Further, there is no way in which subsequent developments of the texts (2) or (3) can induce

the English speaking reader to change his mind about the inferences (a, b) of Table 1.

Inferences (a) and (b) were the same for texts (2) and (3). There are inferences of the same general kind in which the two texts differ. In particular, in text (2) the Orioles' shortstop is the batter while in text (3) the Orioles' shortstop is not the batter. Table 2 adds these inferences to those of Table 1.

	Text (2)		Text (3)
(a)	the ball = it	(a)	the ball = it
(b)	who = the first baseman	(b)	who = the first baseman
(c)	the Orioles' shortstop = the batter	(c)	the Orioles' shortstop ≠ the batter

Table 2

But now we note that these new inferences (c) are in part dependent on schematic knowledge of the game of baseball. In particular, they are dependent on the knowledge that shortstop is the name of a fielding position, that batters may be referred to by their fielding positions, that bunting is something that batters (not fielders) do, that throwing is something that fielders (not batters) do, that the Orioles is the name (or could be the name) of a baseball team, and so on (the list is not complete). The reader who has absolutely no knowledge of baseball will notice that he or she was unable to derive inferences (c).[4]

It is not impossible to imagine a continuation of the texts (2) and (3) that could vitiate inferences (c). We can imagine a game that is generally like baseball except the offensive player is not equipped with a bat but rather attempts to catch the pitch

and throw it out into the playing field. The defensive players, on the other hand, are equipped with bats and are required to hit or bunt the ball to the basemen to put the batters out. A skilled writer might be able to build a text that took (2) or (3) as the initial fragment and develop it along these lines so that the reader would ultimately reject his or her initial inference of type (c) and arrive at a final envisionment of a game of the sort just suggested. A text that began with fragments (2) or (3) and caused us to overthrow inferences (c) might seem far-fetched, perhaps some kind of literary tour de force. Nevertheless such a text is possible in English. In this important respect inferences of type (c) differ from those like (a) and (b), although one suspects that with regard to actual texts and actual readers the analytical distinction is unlikely to make much practical difference. Nonetheless, in thus distinguishing inferences (a, b) from (c), we have detected what we may call two levels of envisionment: the (a, b) level being absolutely ordained by the text and the (c) level being strongly -- though not absolutely -- determined by the text. The point, which we shall now pursue, is that items of the envisionment may be warranted by the text to varying degrees and therefore to varying degrees cancellable by further development of the text without discomfort to the ideal reader.

Consider now the inferences, obvious enough to those fairly familiar with baseball, that in text (2) the first baseman is not an Oriole while in text (3) the first baseman is an Oriole. These inferences require grammatical knowledge and the knowledge of baseball (and the lexicon of baseball) of the kinds already considered. In addition they are based on the world knowledge that the players on one team of a given game normally belong to one ball club while the players

on the other team belong to another ball club. That is, we normally think of baseball games which are known to involve players on regularly organized, league ball clubs as being regular league games, in which the players represent the teams of which they are regular members. There are, however, regularly occurring, if less frequent, events in American culture in which these conditions do not hold, for example, all-star games. The rosters of all-star teams are necessarily composed of players from a variety of teams. If the game being described in text (2) or (3) turns out to be an all-star game, then the text provides no warrant for the inferences just mentioned, which are given as (d) in Table (3).[5]

	Text (2)		Text (3)
(a)	the ball = it	(a)	the ball = it
(b)	who = the first baseman	(b)	who = the first baseman
(c)	the Orioles' shortstop = the batter	(c)	the Orioles' shortstop ≠ the batter
(d)	the first baseman ≠ an Oriole	(d)	the first baseman = an Oriole

Table 3

Inferences (d), although perfectly natural to draw from these texts are, as we have just seen, relatively easily suspendable under culturally non-farfetched, contextual assumptions; this is the same as saying that they are easily suspendable under relatively non-farfetched continuations of the same texts. We want to say, therefore that inferences of types (a, b), (c) and (d) belong to increasingly 'higher'

levels of envisionment, where the lower the level the more direct and absolute the textual warrant for the inference and the higher the level of envisionment the more contingent and revisable the inference is. Roughly speaking, if a text requires revision of something in the envisionment at the lowest level we say that the text is inconsistent, self-contradictory, or incoherent; if the text requires revision of something at a 'medium' level of envisionment we are inclined to say that the text is surprising; if we find that the text requires revision of something we have at a very high level of envisionment we are inclined to attribute no particular property to the text. Rather, we may register surprise that we as readers have populated our envisionment of a text with some item that lacks textual warrant.

For example, suppose the reader envisioned the Orioles' shortstop as right handed. This would be reasonable as, not only are the majority of people right handed, but the exigencies of baseball play are such that, very few if any professional shortstops are left handed.[6] The reader who envisioned the Orioles' shortstop as right handed would have a right to be mildly surprised if it turned out later in the text that he wasn't, and the entitlement to surprise would increase with the length of text that intervened between the introduction of the shortstop and the revelation of his left-handedness.

Compared to the (d) inferences about which team the first baseman is on, this inference about handedness is less directly warranted by the text and more by general background knowledge of baseball; it also seems less likely to figure in coherence relations of the kind we will discuss under the parsimony principle below. For these reasons, we may with fair confidence add it to our table as representing a still more tenuous level of inference.

	Text (2)	Text (3)
(a)	the ball = it	(a) the ball = it
(b)	who = the first baseman	(b) who = the first baseman
(c)	the Orioles' shortstop = the batter	(c) the Orioles' shortstop ≠ the batter
(d)	the first baseman ≠ an Oriole	(d) the first baseman = an Oriole
(e)	The Orioles' shortstop is right handed.	(e) the Orioles' shortstop is right handed.

Table 4

At a level of inference or imagination more tenuous yet, suppose some reader forms an image of the shortstop with a mustache. If it turns out later that the shortstop is clean shaven, any reader who experiences surprise is not, we think, entitled to account for this experience by saying that the text is surprising. Hence we do not add this kind of inference to the ideal reader's envisionment at any level.

The ideal reader will thus make inferences regarding the content world of the text at a variety of levels. As we have noted, the reader who cannot or will not read between the lines is not a good reader, and the reader who is unable or unwilling to distinguish what he or she has been told from what he or she has inferred or imagined is also less than ideal. The latter ability, to distinguish what is directly warranted by the text from that which is less directly warranted (even to the point of what is not warranted at all by the text but is merely contributed idiosyncratically to the envisionment), is most obviously necessary to the reader of informational texts, or, more exactly, to the reader who approaches a text for the purpose of assimilating the information it contains.[7]

We have seen that the reader of informational texts -- or the reader of texts in the informational mode -- must keep straight the levels of his or her envisionment. The informational mode is the mode of reading taught initially in our grammar schools, and some educational critics would claim that it is the only mode of reading ever taught in our society, even in college literature courses. It might be thought, and has indeed been argued, that this kind of self-awareness in the reader -- this ability of the ideal reader to reflect upon its own processing and recognize those parts of the envisionment that derive from distinct sorts of processing routines -- is relevant only, or primarily, to the informational mode of reading. In the coming section I argue that reflexive self-awareness of the reader's internal processing is equally important in expressive-esthetic kinds of reading, and that literary and rhetorical effects, including humor, may require just such awareness on the part of the reader of his or her own processing of the text.

The Trusting Ideal Reader

We have seen that in order to study the interaction of young readers and the texts they encounter on tests of reading comprehension we must postulate an ideal reader with the degree of sophistication necessary to distinguish appropriate levels of the envisionment. There is another kind of sophistication possessed by real, mature readers that it is for the most part not necessary to impute to the ideal reader with which beginning readers are compared. This impoverished kind of ideal reader is one that lacks the capacity to reflect upon the history of its own processing of a text. The TRUSTING ideal reader is literal minded. If it form a hypothesis and then encounters evidence that leads it to reject that hypothesis, it doesn't have to ask itself second order questions such as, "Why would the author have put in the information that led me to the first hypothesis if he or she were only to give me conflicting information later? Is there some kind of literary trick going on here?" The trusting ideal reader is artless in the sense that it does not produce secondary interpretations of a text by reflecting upon the processing it has done in producing the primary interpretation. The trusting ideal reader does not construct an image of the author of the text it is processing and does not reflect upon this author's possible motivations.

The best illustration I can think of for the concept of the trusting ideal reader involves a text intended, not primarily to be read, but rather to be heard: a bit of doggerel that my peers and I in junior high school found amusing:

(7) Of all the things I'd rather be,
I'd rather be a bass.

I'd climb up to the top of trees

And slide down on my hands and knees.

Part of the intended humorous effect of the text depends on the whimsical image of fish climbing trees; we will leave this matter for the moment and return to it later. For the rest, the addressee who gets the point and intended humor of the text incorporates the trusting addressee and reflects upon the results of this very trustingness in arriving at the intended rhetorical effect.

From the point of view of the trusting addressee the analysis of (7) goes in part as follows. At the end of the third line, one has noted that the rhyme scheme is so far a b c (or perhaps a b a, depending on whether trees is heard as an imperfect rhyme for be), and with regard to meter that all three lines are perfectly iambic, having four, three, and four feet respectively. The pattern of rhyme and meter so far perceived is

(8) ⏑ – ⏑ – ⏑ – ⏑ –

⏑ – ⏑ – ⏑ [bǽ s]

⏑ – ⏑ – ⏑ – ⏑ –

One therefore contracts the expectation that the next (i.e., the fourth) line will (1) end in a stressed syllable concluding with [æ s] to rhyme with bass and (2) contain three iambic feet, completing the pattern given in (7) which is the basic ballad form, an extremely familiar genre in our culture.[8]

(9) ⏑ – ⏑ – ⏑ – ⏑ –

⏑ – ⏑ – ⏑́ X

⏑ – ⏑ – ⏑ – ⏑ –

⏑ – ⏑ – ⏑́ X

At the end of the word my in the fourth line -- with one syllable left to go -- these hypotheses are still working perfectly. We lack only a single stressed syllable to complete the ballad form (4 3 4 3 iambs), and there is a monosyllabic word which rhymes with bass, which names a part of the body one can slide on, thereby making as much semantic sense as one may hope for in a text containing the whimsy of bass climbing trees, and which, above all, provides the satisfaction of completing the pattern of rhyme, meter, and meaning with a tabooed word. For these reasons the trusting ideal hearer expects that word at this point.

When line 4 continues hands and knees the trusting addressee cancels his expectation without further thought and notes with simple satisfaction that the poem has been successfully concluded with a rhyme scheme a b c c and a matching meter of 4, 3, 4, 4 iambs. The trusting ideal addressee does not reflect upon the fact that his expectations have been violated, he merely revises them. He therefore doesn't appreciate that a trick has been played on him: that he has been led to contract an expectation for a tabooed word on the basis of the rhythmic, metric, and semantic pattern, but that that pattern has been completed satisfactorily without the use of the tabooed word. The sophisticated addressee, on the other hand appreciates that a joke has been played on her in that she has been convicted in her own mind of 'dirty-mindedness' for having supplied the tabooed word before it was produced, only to find that the poem was successfully completed without it.

What we want to note here is that the sophisticated ideal reader incorporates the trusting ideal reader. The sophisticated ideal reader gets the joke by virtue of noticing how the trusting reader inside of

her or him has been led up the garden path. It seems a useful working hypothesis that many if not all texts that require something more than our trusting ideal reader to appreciate their full rhetorical effect require as one component of their sophisticated ideal reader a trusting ideal reader. In so far as this hypothesis is correct, any theory that aims to explain the processes used by sophisticated readers in interpreting literary texts that rely on complex rhetorical strategies will need to include a theory of a trusting ideal reader. We therefore have grounds to hope that what we learn about the reading process by positing a trusting ideal reader who is competent to read very simple texts may be of use in understanding the workings of sophisticated readers who are able to interpret more complex texts successfully.

Linda Coleman, of our project, has drawn attention to examples, such as the following, which illustrate simultaneously both the levels-of-envisionment property and the trusting-versus-sophisticated-ideal-reader property. In a science fiction novel, Ursula Le Guin introduces a character as the Terran ambassador (i.e. the ambassador from Earth). Some time later, after this character has played a minor but not insignificant role in the proceedings for a while, we are allowed to learn that the Terran ambassador is a black woman. This comes as a shock to those readers, probably all, who have envisioned the ambassador from Earth as a white man. But, there is a double shock because the reader is shocked at being shocked. The initial shock is connected with feelings that the author is skating dangerously close to incoherence by allowing such a strongly warranted inference to stand this long before correction. The second shock comes with the rapid realization that the very strength of the inference is based on the strength of racial and sexual stereotypes

that the reader of this kind of literature probably thinks that he or she doesn't have in the first place. And then the sophisticated reader envisions the rhetorical-political purposes of the author, who has forseen all this processing going on in her readers' minds and intended it to happen just this way. By having kept straight his or her levels of envisionment the reader is able to reject the image of a white man in the envisionment and substitute that of a black woman without imaging a contradiction. Further reflections on his or her own stereotyped behavior, racism, sexism, or on images of a clever author, non-didactically pointing these out, belong to the sophisticated readers' reflection on the trusting reader inside her or him, who was able to effect the substitution 'black woman' for 'white man' but appreciate nothing of its significance.

The Parsimony Principle

In the first section of this paper, where we considered the necessity of keeping track of the different levels of certainty of the envisionment, we perforce considered the ideal reader's ability to read between the lines of a text, that is, to draw inferences that are not directly warranted by the grammar and lexicon of the text but that are nonetheless necessary to derive a coherent envisionment. For example, from a text such as (10), the reader who fails to draw inferences such as (11 a-g) will not be an ideal reader.

(10) One day a chef went to Fisherman's Wharf and bought some fish from a fisherman.

(11) a The chef will cook the fish at his restaurant.[9] He will not, for example, take it home to his wife.

b The fisherman caught the fish. For example, he is not

selling it for an electrician friend who happened to buy too much fish at the supermarket.

c The fisherman is a commercial fisherman, and he caught the fish in order to sell it. He is not, for example, a lucky sports fisherman who just happened to be opportuned by the chef.

d The money the chef gave the fisherman was not his own money but the restaurant's.

e The purpose of the chef's visit to Fisherman's Wharf was to buy fish, not, for example, to visit the Wax Museum.

f The transaction took place at Fisherman's Wharf. Compare: "One windy day Charlotte went to the hairdresser's and bought a bandana from a street vendor."

Each of the inferences of (11) is of the between-the-lines type, that is, it is not absolutely warranted by the grammar and lexicon of the text. In each case the contrasting possibility mentioned is not what we immediately infer but is something that could turn out to be what the author had in mind. Inferences (11 a-g) are like the middle level inferences (c, d, and perhaps e) in our baseball example (2, 3); the ideal reader of this text would draw these inferences but would also remember them as, in varying degrees, subject to suspension by possible later developments. If they are suspended by later developments, the reader will be entitled to a reaction of surprise; the longer the delay, the greater the surprise.[10]

So far we have discussed things that happen to the ideal reader's inferences after they are drawn, but we have not talked about how they are drawn in the first place. A full theory of the ideal reader will be concerned with the latter quite as much as the former. Although a complete theory of the ideal reader, one which predicts the actual inferences drawn from a text without making use of the human judgement of the analyst, seems distant from the present state of the art, perhaps there are a few steps we can take in this direction. Let us proceed inductively in the analysis of text (10) and the inferences (11) to which it gives rise, first discovering the background schemata that seem necessary and then investigating how these schemata are employed in reaching these inferences.

First, the word chef evokes a RESTAURANT schema. We have already noted (Footnote 9) that chefs may not always occur in restaurants and we may at times wish to speak of a CHEF schema per se. Hence: in semantic memory schemata may intersect each other in a variety of ways: the CHEF schema is itself an (optional) element of the RESTAURANT schema and is also an instance of an OCCUPATIONAL schema.

There is also a FISHING schema evoked by our text. In fishing there is a person who attempts to catch fish, sometimes succeeds, may own fishing equipment, may ride in a boat, etc. There are two principal kinds of fishermen: commercial fishermen and sports fishermen. One way to give theoretical recognition to this observation is to say that the fishing schema has two sub-schemata, sport and commercial. A contrasting way is to say that the fishing schema itself instantiates two distinct families of schemata: OCCUPATIONS and AVOCATIONS. A

formal theory of semantic memory, one that specifies not only its constituent schemata but also the relations amongst them, would presumably have to take a position on issues such as this. For present purposes we are content to discover some of the elements of such a theory without specifying their interrelations.

The well known COMMERCIAL EVENT schema (Fillmore 1977) is of course involved in text (10). We have made use in a casual way of the concept family of related schemata. In the present case, we may note that for all their many dissimilarities the FISHING and COMMERCIAL EVENT schemata have something in common. They both involve purposeful actions and hence each involves at least one goal.

There is also a schematization of the time relations of actions involved in this, as in any text. (The temporal schematization of (10) is discussed in Foot note 10.)

We would like now to say something about how these schemata are employed in drawing the between-the-lines sorts of inferences listed in (11). As a preliminary, let us consider how schemata are involved in drawing ground-level, unsuspendable inferences from our text (10). One such unsuspendable inferences is:

(12) The person who went to Fisherman's Wharf is the same as the one who bought the fish.

We can put this more long-windedly but in a way that better reveals our theoretical preconceptions, as follows: the grammar of (10) guarantees that the entity who fills the TRAVELER slot of the JOURNEY schema evoked by the first clause is the same entity as the one that fills the BUYER slot of the COMMERCIAL EVENT schema evoked by the second clause; the TRAVELER participant in the JOURNEY scenario that forms the first part

of our envisionment and the BUYER participant in the COMMERCIAL EVENT scenario that forms the second part of our envisionment are the same. We are talking here of schemata as structures in semantic memory that are employed on particular occasions to build the scenarios that constitute an envisionment. We speak of schemata as containing slots and of the things that fill these slots in a particular scenario as participants. Often, the slots in schemata that are filled by participants in scenarios are matched by noun phrases in the grammatical structure and the question whether two participants in different scenarios are the same often corresponds to the question whether two noun phrases are coreferential. We see thus that the main thing happening in inference (12) is that a single participant, the chef, connects two scenarios (going to Fisherman's Wharf, buying a fish) by filling distinct slots (traveler, buyer) in the schemata (JOURNEY, COMMERCIAL EVENT) that underlie the two scenarios. The two little scenarios are thereby joined into one larger scenario, giving the text coherence.

We have already noticed that, unlike the inferences of (11) with which we are primarily concerned here, the inference of (12) is forced by the grammar. If we consider a text of parallel grammar but which lacks any intuitive semantic support for an inference comparable to (12), we see that such an inference (14) is forced anyway. That is,

(13) One day a chef went to Fisherman's Wharf and sprained his ankle.

(14) The person who went to Fisherman's Wharf is the one who sprained his ankle.

unlike the inferences (11), the grammar forces the inference of (12) from (10) as it does the inference of (14) from (13).

This aspect of sentence semantics seems directly related to a principle of text semantics that we call the 'parsimony principle'. The parsimony principle is this: whenever it is possible to link two separate scenarios into a single larger scenario by imagining them as sharing a common participant, the ideal reader does so. Let us turn directly to our examples (11) to see how this principle works out in practice. In (15 a-g) the inferences (11 a-g) are repeated and each is followed by an explanation of how it arises from (10) via the parsimony principle.

(15) a The chef will cook the fish at his restaurant. The food participant in the restaurant scenario and the goods participant in the commercial event scenario are the same.

b The fisherman caught the fish. The fish participant in the fishing scenario is the same as the goods participant in the commercial event scenario.

c The fisherman caught the fish in order to sell it and he is a commercial fisherman. The entire commercial event scenario is the goal participant of the fishing scenario. In the commercial fishing schema (or sub-schema) the fisherman is also a fish seller. If the commercial fishing schema is chosen, then the seller slot of the commercial fishing scenario can be occupied by the same participant as the seller slot of the commercial event scenario. (This requires something beyond the parsimony principle as baldly stated; we will come back to this point.)

d The money the chef gave the fisherman was not his own but

the restaurant's. The restaurant schema has a slot for money, since restaurants are profit seeking enterprises. The money slot in the restaurant schema can be made to share a participant with the money slot on the commercial event schema.

e The goal of the chef's visit to Fisherman's Wharf was to buy fish. The entire commercial event scenario can serve as the goal participant of the journey scenario.

f The transaction took place at Fisherman's Wharf. The destination participant of the journey scenario is inferred to be the same as the (optional) location slot of the commercial event scenario and the (optional) home-port slot of the fishing scenario.

These equations of participants across scenarios are displayed graphically in Figure 2 by the vertical arrows with single shafts. (Only the double-shafted arrows correspond to equations directly mandated by the language of the text and therefore requiring no potentially suspendable inferences.) Words in *italics* indicate words of the actual text. The words following these in CAPITALS are the names of schemata that are evoked by the words of the text. Following the name of a schema is a series of names of slots in underlined lower case. Following the name of each slot, after a colon, is a word in ordinary print intended to indicate the identity of the participant of that slot in the particular scenario created by this text. For example, the arrow marked 'a' indicates that the goods participant of the COMMERCIAL EVENT scenario is the same as the food participant of the RESTAURANT scenario; the fact that this arrow is single shafted indicates that this is a suspendable inference and the fact that it is marked 'a'

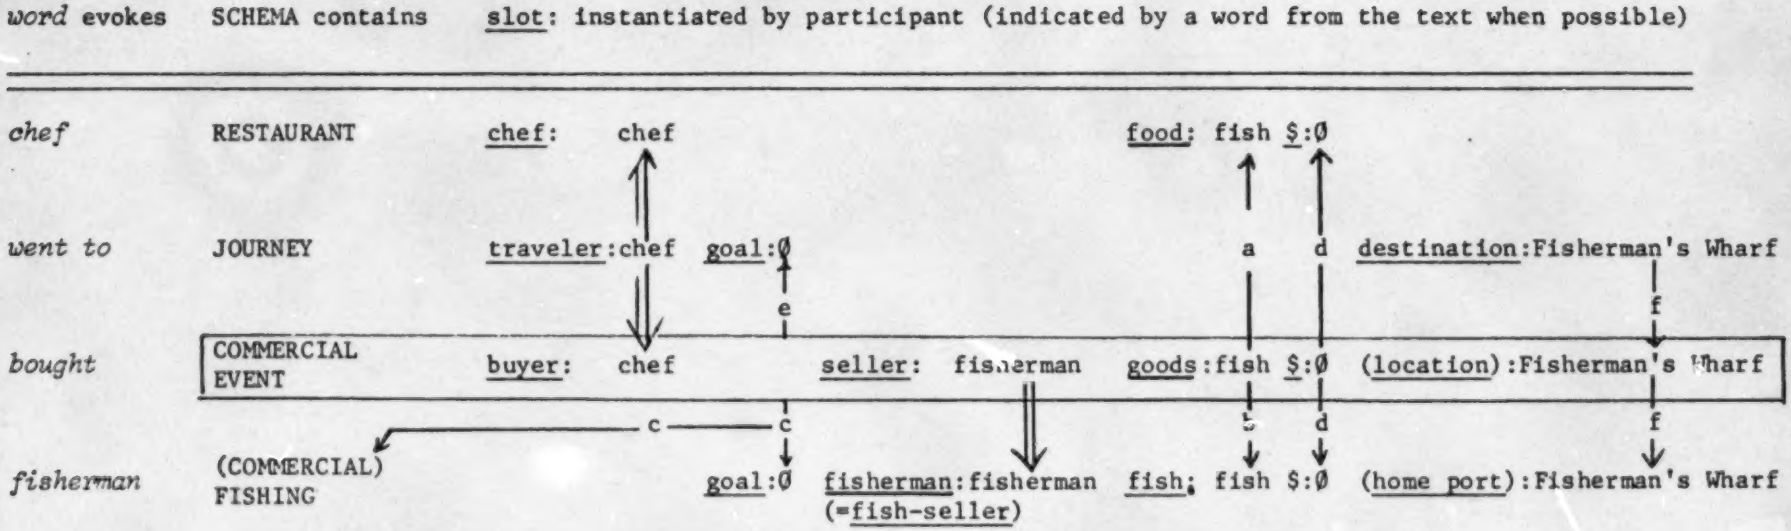

Schemata are evoked by words of the text. Schemata contain slots. Scenarios are instantiations of schemata in which the slots are filled by particular envisioned participants. The grammar and lexicon of the text and, particularly, the parsimony principle, create equations among participants in different scenarios, binding these individual scenarios into the larger scenario that is the envisionment.

Figure 2.

indicates that this inference corresponds to the inference labelled (11)a and (15)a in the text. The list of slots appearing for a given schema is not intended to represent a comprehensive analysis of that schema but only to portray those slots whose participants are identified with the participants of scenarios based on other schemata in the envisionment of this particular text (by the ideal reader). The inferences indicated by single arrows a - f on the Figure are instances of the parsimony principle in action.

The examples of the parsimony principle so far considered have had the following structure: (1) words of the text evoke schemata, (2) these schemata contain slots that need to be filled by participants in the text-specific scenarios that instantiate them, (3) the parsimony principle enjoins the ideal reader to keep the number of distinct participants to a minimum, hence to equate participants in distinct scenarios whenever possible, (4) this results in the entwining of the various small scenarios, creating out of them one large, reticulate scenario. I would suggest that this reticulation of the overall scenario (i.e., of the envisionment) is an important part of our intuition of textual coherence.

We have concentrated on examples in which the words of the text evoke schemata and the parsimony principle then goes to work on these. But we have also noted in passing certain examples where the parsimony principle does even more: it helps to direct the initial choice of relevant schemata or sub-schemata. One such example involves the choice of the commercial fishing subschema over the sport fishing subschema. We note that if we chose the commercial fishing subschema, then we might get a tighter schematic fit between two participants who

must be matched anyway on account of strict grammatical schemata. Recall that the grammar ordains that the fisherman participant of the fishing scenario and the seller participant of the commercial event scenario be the same. By choosing the commercial fishing sub-schema, in which fishermen are fish sellers, we achieve a matching in semantic schematization of participants who are anyway identified with each other by the grammar of the text. Here we see that the parsimony principle not only matches participants for schemata that are already evoked but leads the ideal reader in his selection of the content schemata themselves. This suggests a sort of metaprinciple for schema selection: select schemata in such a way as to give the parsimony principle the widest possible scope of operation. We might dub this meta-principle the parsimony promotion principle. Much more empirical work is obviously needed to delineate the details of the operation of the closely related parsimony and parsimony promotion principles, but even the very preliminary examples discussed here suggest that these principles are real and play an important role in both ideal readers' and competent real readers' interpretation of texts.

Footnotes

1. Co-principal investigators on the project, supported by NIE Grant No , are Charles Fillmore and myself. Reading specialist Judith Langer has participated, as have graduate student assistants Robert Aronowitz, Karen Carroll, Linda Coleman, Katherine Kovacic, Thomas Larsen, and Mary Catherine O'Connor. Welcome advice has been received from Mary Sue Amon, Kjell Madsen, John R. Ross, and Patrizia Violi.

2. The texts with which our project is primarily concerned are the items on tests of reading comprehension given to third and fifth graders. An alarming number of these texts are deficient in the sense of requiring an ideal reader that corresponds either to no real reader (no matter how proficient) or to a real reader that it is unreasonable to expect a third or fifth grader to approximate. Some examples of such texts are given by Fillmore (1981), who also discusses the role that the concept 'ideal reader' plays in our empirical procedures for studying the interaction between third and fifth graders and these texts.

3. Reading is probably easier to model in this way than hearing is because reading lacks the interactive phenomena, such as negotiation of interpretive context, back channel signalling, and the like, that are the bread and butter of the conversational analyst. Nevertheless, reading is hard enough. In the remainder of this paper I talk about reading, but to the extent the model we are developing is correct, it should be useful in the understanding of the text semantics of spoken language as well. Written language also differs from spoken language in other ways, such as the imperfect substitution of punctuation for intonation, stress, and prosody.

4. Readers who doubt that inferences (c) require some knowledge of baseball should consult an acquaintance who lacks all such knowledge. To make the point I am making here, it has been necessary to select an example from a domain about which knowledge is widespread but not ubiquitous among English speakers. Those who lack knowledge of baseball at the requisite level will have either to take my word for it or consult someone who has such knowledge. No affront to either knowers or non-knowers of baseball is intended.

5. There are other plausible continuations of texts (2) and (3) in which inferences (d) need not -- in some cases cannot -- hold; examples include sandlot games and intra-squad games.

6. Left handedness is common at other positions; e.g., pitcher and first baseman.

7. Some texts -- for example, assembly instructions -- seem inherently informational while others -- for example, poems -- seem inherently intended to arouse aesthetic experiences rather than to impart objective information. It is, however, possible and perfectly normal to approach texts originally intended to be taken one way in another. Such is the case, for example, when the literature student reads a play or poem to learn its content and structure rather than to experience its aesthetic values. There is, of course, much more to the subject of what functions a text may fulfill than is suggested by the simple contrast 'impart objective information' vs. 'arouse aesthetic experience'.

8. One of many familiar quatrains that might be cited is from Wilde's Ballad of Reading Gaol:

8. (continued)

For each man kills the thing he loves;
Let this all be heard.
Some do it with a bitter look.
Some with a flattering word.

Note the rhyme of heard and word

9. For simplicity we assume here that chefs are employed only in restaurants, though it is of course true that chefs are sometimes employed in the homes of the rich. Another assumption we have made tacitly, that has perhaps been noticed by some readers, but probably not many, is that the chef is a man. Such inferences from stereotypes are a major target of various kinds of language reform; no doubt something of this reformist spirit motivated the novelist Le Guin when she constructed the Terran ambassador text mentioned above.

10. Inferences (11 a-g) by no mean exhaust what the ideal reader should extract from this text. In particular, the ideal reader will construct several little histories for the various participants in this story and know the temporal points of their relation to each other. Thus the money that paid for the fish probably originally came from customers in the restaurant who paid for food there and may well end up being paid out for bait or gasoline or other fishing supplies. The catching of the fish by the fisherman took place before the commercial event, while the cooking by the chef will take place subsequently, and the eating by the customers of the restaurant later still. Some, but not all, of the relevant historical understanding is represented diagrammatically in

10. (continued)

in the following picture (Figure 1). At $time_1$ the chef has not got hold of this particular money and the fisherman is not yet in possession of the fish. Sometime later the money and the chef have become associated and the fish and the fisherman have become associated; this is represented by the picture at $time_2$. The commercial event takes place at $time_3$, briefly bringing all four participants into spatio-temporal contiguity. Following this, at $time_4$, the four participants divide into novel pairings, fish with chef and money with fisherman. Eventually, $time_5$, these pairings dissolve as the fish is eaten and the money spent or invested.

The relations depicted at times 2, 3, and 4 follow from our knowledge of the commercial event schema itself, while those at earlier and later points in the history, represented by times 1 and 5, depend on our schematic knowledge of the kinds of participants involved in the particular commercial event portrayed in this text. For example, with reference to $time_1$, it follows from our knowledge of FISHING (not of commercial events) that the fish and the fisherman were not always associated. If we consider a commercial event with different kinds of participants, we are not led to imagine an earlier time when the seller and the goods have not yet become associated. For example in O.Henry's Gift of the Magi the sale of the young matron's beautiful hair does not invite us to imagine an earlier time when hair and matron enjoyed separate existences. Similarly, the ideal reader who encounters a commercial event in which someone buys an artificial limb or a pacemaker will not forsee a future time when the purchaser and his acquisition have become separated.

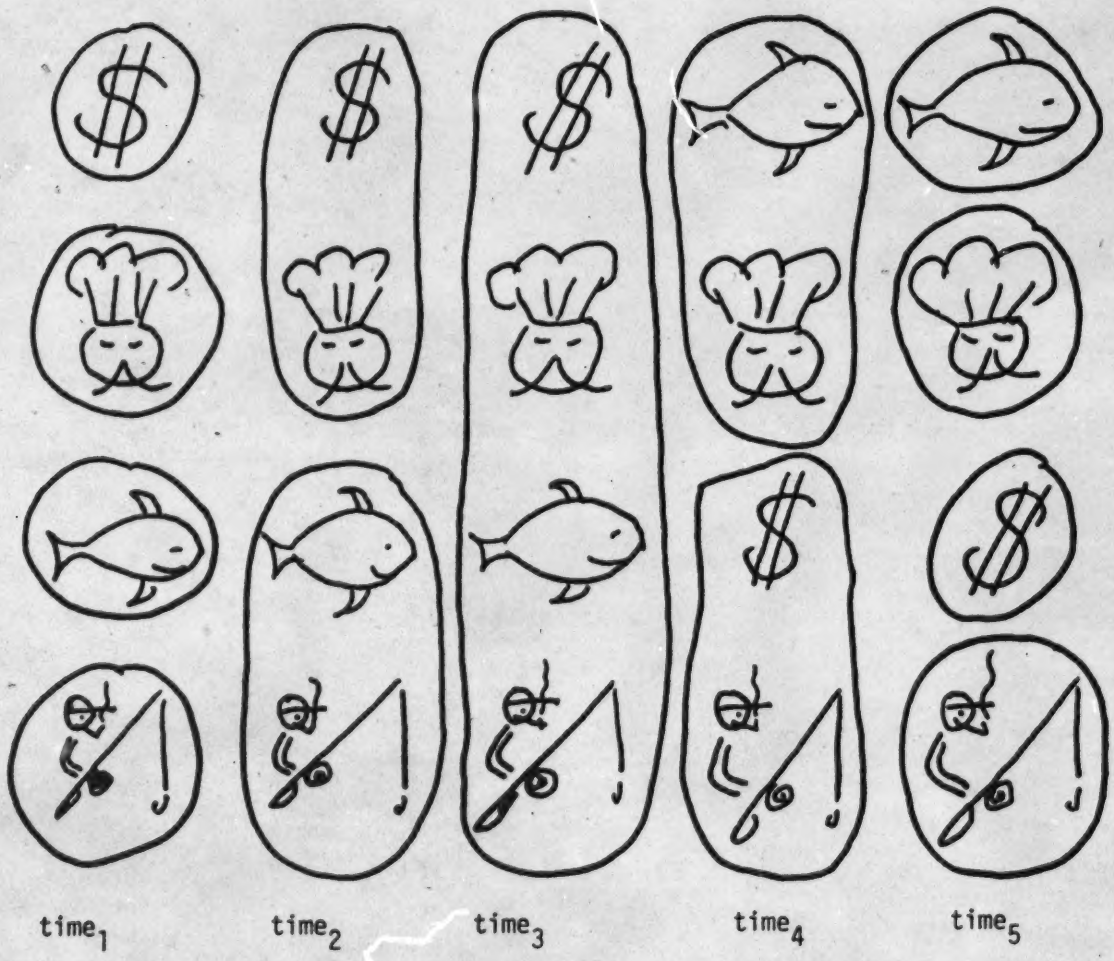

Figure 1

10. (continued)

Of the relations depicted at times 1 and 5 in Figure 1, the separation of fish and fisherman at $time_1$ seems to have the most direct warrant in the text, in the form of a fishing schema which insists that fishermen catch fish they didn't previously have. (An activity, otherwise like fishing, that takes place in a bucket or a bathtub is probably not real fishing and surely not prototypical fishing.) In the actual analysis of any text, one always reaches a point at which the inferences that different competent analysts attribute to the ideal reader diverge. For example, some might wish to say that the separateness of the money and the chef at $time_1$ is something that the reader is supposed to get out of this text, while others will say that this is a plausible imagining, consistent with the text, but not something that would cause us to say that a reader had missed something if he or she failed to come to it. The methodological point is that for any text there are many things that all analysts agree in attributing to the ideal reader and also things about which they will disagree. These latter define those levels of envisionment where what is ordained by the text shades imperceptibly into what is permitted by the text.

REFERENCES CITED

Fillmore, Charles J.

1981 Ideal readers and real readers. Proceedings of the 32nd Georgetown Roundtable on Languages and Linguistics. Washington, D.C. Georgetown University Press.

1977 Topics in lexical semantics. In R. Cole (Ed.) Current Issues in Linguistic Theory. Bloomington, Ind. Indiana University Press.

Fillmore, Charles J. and Paul Kay

1980 Progress Report: Text Semantic Analysis of Reading Comprehension Tests. Typescript 80 pp.